

Active learning of potential-energy surfaces of weakly bound complexes with regression-tree ensembles

Cite as: J. Chem. Phys. **155**, 144109 (2021); <https://doi.org/10.1063/5.0057051>

Submitted: 17 May 2021 • Accepted: 21 September 2021 • Published Online: 13 October 2021

 Yahya Saleh,  Vishnu Sanjay, Armin Iske, et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Nearly reducible finite Markov chains: Theory and algorithms](#)

The Journal of Chemical Physics **155**, 140901 (2021); <https://doi.org/10.1063/5.0060978>

[Spherical potential functional theory](#)


The Journal of Chemical Physics **155**, 144108 (2021); <https://doi.org/10.1063/5.0065942>

[Gaussian process models of potential energy surfaces with boundary optimization](#)

The Journal of Chemical Physics **155**, 144106 (2021); <https://doi.org/10.1063/5.0063534>

Challenge us.

What are your needs for
periodic signal detection?



Zurich
Instruments



Active learning of potential-energy surfaces of weakly bound complexes with regression-tree ensembles

Cite as: J. Chem. Phys. 155, 144109 (2021); doi: 10.1063/5.0057051

Submitted: 17 May 2021 • Accepted: 21 September 2021 •

Published Online: 13 October 2021



Yahya Saleh,^{1,2,a)} Vishnu Sanjay,^{1,2,3} Armin Iske,² Andrey Yachmenev,^{1,3,a)} and Jochen Küpper^{1,3,4,5}

AFFILIATIONS

¹ Center for Free-Electron Laser Science CFEL, Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany

² Department of Mathematics, Universität Hamburg, Bundesstraße 55, 20146 Hamburg, Germany

³ Center for Ultrafast Imaging, Universität Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany

⁴ Department of Chemistry, Universität Hamburg, Martin-Luther-King-Platz 6, 20146 Hamburg, Germany

⁵ Department of Physics, Universität Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany

^{a)} Authors to whom correspondence should be addressed: yahya.saleh@cfel.de and andrey.yachmenev@cfel.de.

URL: <https://www.controlled-molecule-imaging.org>

ABSTRACT

Several pool-based active learning (AL) algorithms were employed to model potential-energy surfaces (PESs) with a minimum number of electronic structure calculations. Theoretical and empirical results suggest that superior strategies can be obtained by sampling molecular structures corresponding to large uncertainties in their predictions while at the same time not deviating much from the true distribution of the data. To model PESs in an AL framework, we propose to use a regression version of stochastic query by forest, a hybrid method that samples points corresponding to large uncertainties while avoiding collecting too many points from sparse regions of space. The algorithm is implemented with decision trees that come with relatively small computational costs. We empirically show that this algorithm requires around half the data to converge to the same accuracy in comparison to the uncertainty-based query-by-committee algorithm. Moreover, the algorithm is fully automatic and does not require any prior knowledge of the PES. Simulations on a 6D PES of pyrrole(H₂O) show that <15 000 configurations are enough to build a PES with a generalization error of 16 cm⁻¹, whereas the final model with around 50 000 configurations has a generalization error of 11 cm⁻¹.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0057051>

I. INTRODUCTION

The ability of molecules to form weakly bound complexes, via van der Waals or hydrogen bonds, is fundamental to many physical, chemical, and almost all biological processes, ranging from the phenomenon of resonances in reactive scattering¹ and roaming in chemical reactions² to biochemical processes in aqueous solution.³ Small microsolvated complexes of aromatic molecules or chromophores are important model systems for studying the interactions established by individual water molecules at specific binding sites.⁴ They offer an appealing way for the quantitative investigations of the effects of hydration on the photoexcitation and

ionization dynamics and the mechanisms of chemical reactions in general.^{5–10}

First principles calculations of potential-energy surfaces (PESs) for weakly bound complexes are challenging and computationally expensive.^{11–14} The presence of small energy differences in weakly bound PES means that, generally, high levels of theory need to be employed to produce correct asymptotic behavior.¹⁵ Furthermore, the landscape of these PESs is complex because of the loosely bound character of intermolecular interactions. Thus, a larger number of grid points is generally required to sample the complete configuration space. Moreover, due to the importance of dynamical electron correlation (dispersion) and its slow basis-set convergence,¹⁶

calculations for the noncovalent long-range parts of the PES are generally more costly than the ones at short-range.

The standard approach for building *ab initio*-based molecular PESs employs a regression procedure: one generates a grid of distinct molecular geometries at which the potential energy is computed, and then a model is fitted to these data by minimizing a loss function, e.g., root-mean-square (rms) error, under a regularization constraint.^{17–19} In recent years, many machine learning (ML) models^{20–24} have been used to fit molecular PESs. The most extensively used models include permutationally invariant polynomials,^{25–29} neural networks (NNs),^{30–36} Gaussian processes (GPs),^{37–43} and other kernel methods.^{44–47}

The quality of PESs strongly depends on the chosen set of molecular geometries sampled. This choice is often based on humans' intervention: the regions that will likely be intensively accessed by the dynamics simulations are sampled more densely. If one's assumptions about the distribution of geometries accessed in dynamics' simulations are mostly correct, the question is what minimal amount of data points is needed to reproduce the PES to a desired accuracy.

Trying to minimize the number of electronic structure calculations and reducing the amount of humans' intervention in an active learning (AL) paradigm⁴⁸ became increasingly popular during the last few years.^{49–61} AL encompasses a variety of iterative algorithms aimed at minimizing the cost of training data acquisition. In pool-based AL, the expert defines a pool of unlabeled data, e.g., molecular geometries, and a policy algorithm queries the energies of geometries that would most significantly minimize the generalization error of the PES once labeled and included in the training data.⁴⁸ The vast majority of AL applications to PESs used uncertainty-based sampling,^{49,51–58} which queries points corresponding to the largest uncertainties in their predicted energies. For probabilistic models such as GPs, the uncertainties can be directly calculated.^{52,54,55,60} For the ML models that do not offer a direct way to compute uncertainties, these can be inferred by training an ensemble of models on the currently labeled training set and selecting the points about which the models disagree the most. This algorithm is called query-by-committee (QBC)⁶² and it was used for uncertainty-based sampling of PESs.^{49–51,57,58}

Uncertainty-based sampling is known to be prone to querying outliers^{48,63} because it does not take into account the statistical distribution of data in the pool and hence partly ignores the user's *a priori* knowledge embedded in the pool. Upper bounds on the generalization error in an AL setting suggest that more optimal AL strategies can be built by sampling points corresponding to large uncertainties while not deviating much from the true distribution of data.^{64,65} In applications to PESs, uncertainty-based algorithms tend to select many molecular geometries from regions sparsely sampled by subsequent dynamics calculations. When all sparse regions of the pool can be clearly identified, e.g., as points with high energy, or if prior knowledge about the minima and saddle points exists, this problem can be solved by introducing a weighting function.^{50,60} In a more general setting, one can combine the uncertainty-based query algorithm with a molecular-dynamics sampler starting from various known critical points of the PES.^{25,51,57}

Another concern in AL is the execution speed and the scaling with the amount of data, as AL is an iterative procedure that

includes training and predicting energies of many geometries in each iteration.

Here, we propose a regression version of the stochastic query by forest (SQBF)⁶⁶ selection algorithm and use it to generate a minimal grid for the intermolecular PES of the weakly bound pyrrole(H_2O) complex.^{9,67} The algorithm selects molecular geometries with large uncertainties in corresponding energy predictions and, at the same time, keeps the distribution of selected geometries similar to that in the pool. The proposed algorithm does not suffer from the shortcomings of uncertainty-based methods and leads to better results. The uncertainty is estimated using an ensemble of regression trees that constitute a random forest regressor (RFR), while the data statistics is partially preserved using a random sampling (RS) procedure. Owing to the greedy optimization algorithm employed in regression trees, this method is computationally fast and scales well with the size of data. In general, the RFR model is not a continuous function and hence not well suited for modeling of PESs. However, due to its high computational efficiency, the RFR is an attractive method for sampling purposes in AL, which can be easily combined with other ML models, e.g., NNs, to produce the final model. Our empirical results show that the SQBF selection scheme combined with a NN model for the fitting of PES requires only 30% of the initial pool data (around 50 000 grid points) to achieve the convergence with a generalization error of about 16 cm^{-1} . In comparison to popular QBC and RS schemes, the SQBF queries about two and three times less data, respectively, to converge to the same accuracy. The SQBF algorithm is general, not restricted to weakly bound complexes, and applicable to larger molecular systems.

II. METHODS

A. A formal look into active learning

We denote the joint distribution of the problem as \mathcal{P} , where $\mathcal{P}(x, E) = p(x)p(E|x)$. In standard supervised ML, one has a dataset of points that are independently sampled and identically distributed according to \mathcal{P} .²³ In an AL paradigm, the dataset is selectively sampled from another data distribution $\mathcal{Q}(x, E)$ that has the same conditional distribution of \mathcal{P} , i.e., $\mathcal{Q}(x, E) = q(x)p(E|x)$.^{64,65} We define $P = \{X_1 \dots X_l\}$ to be the set of all distinct molecular geometries in the pool and assume that $P \sim p(x)$. We define $S^{(0)} = \{(X_1, E_1) \dots (X_m, E_m)\}$ to be the initial labeled subset of P , where $m \ll l$. Molecular geometries in the initially labeled set are denoted by $S_X^{(0)} := \{X_1 \dots X_m\}$ so that $P^U = P \setminus S_X^{(0)}$ determines unlabeled geometries in the pool. The steps of a pool-based AL algorithm in a batch-mode⁴⁸ are summarized in Algorithm 1. To decide whether or not the performance is satisfactory, one needs to choose a stopping criterion. We used the root-mean-square (rms) error on a test dataset. We refer to an algorithm that performs step (a) as a policy algorithm.

A simple example of a policy algorithm is uniform RS from the pool. Such an algorithm is representativeness based, as it samples more data from dense regions. In the present case, RS queries more molecular geometries with energies close to the dissociation limit, which is due to the distribution of data in the pool.

An example of an uncertainty-based algorithm is QBC. In this algorithm, one trains an ensemble of efficient and diverse learners on the currently labeled dataset. To decide whether the inclusion

ALGORITHM 1. Basic steps of a generic pool-based active learning strategy. In each active learning iteration, B data points are chosen from the pool, labeled, and added to the training data.

Fix batch size $B, t = 1$;
Input: Pool of unlabeled data P^U , an initial labeled data $S^{(0)}$
while *performance is unsatisfactory* **do**
 a) Select B elements from P^U to obtain $S_X^{(t)}$
 b) Label this set to obtain $S^{(t)}$
 c) Set $P^U = P^U \setminus S_X^{(t)}, S^{(t)} = S^{(t-1)} \cup S^{(t)}$
 d) $t = t + 1$
end

of a certain point would improve the performance of the regressor, these learners are asked to predict the energies of this point. If they produce very different values, we conclude that adding this point to the training set will improve the performance and query the user about the corresponding energy. The algorithm is formalized in Algorithm 2. We note that diversity of learners is crucial in this algorithm. If the learners are not diverse, their predictions for a certain geometry would be almost the same and hence one would not be able to infer the uncertainty. Practically, the diversity of learners is introduced through random perturbations to the learning process.

Uncertainty-based algorithms do not take the statistical properties of data into account when querying points, which is a serious drawback when the distribution of molecular geometries in the grid is nonuniform. We empirically show that QBC tends to sample points from sparse regions to the detriment of high-density parts. To correct this behavior, one can construct a probability density function from the QBC-estimated uncertainties. Then, querying grid points is performed through random sampling according to this density function. Algorithm 3 formalizes this idea. In contrast to QBC, points with small uncertainties, such as points close to the asymptotic limit, may still be queried if they fall in high-density regions. In other words, Algorithm 3 respects the statistical information in the pool that is defined *a priori* by the expert. Note that accounting for the statistical information in the pool using QBC can also be performed by considering only a few unlabeled data points sampled independently from the input distribution as candidates to

ALGORITHM 2. Basic steps of a query-by-committee algorithm for regression problems. n learners are trained on the labeled datasets and asked to make predictions on the whole unlabeled dataset. The B elements chosen to be labeled are those that maximize the standard deviation of the prediction among the n learners.

Fix the number of learners n in the ensemble.
Result: B elements from P^U
Input: $P^U, S^{(t)}, B$
 a) Train an ensemble $\{T_i\}_{i=1}^n$ of models on data $S^{(t)}$.
 b) Compute predictions $\hat{E}_i = T_i(X) \forall X \in P^U, \forall i$
 c) Compute the community disagreement
 $q(X) = \text{std}(\hat{E}_i) \forall X \in P^U$
 d) Take B elements from the unlabeled data that have the highest $q(X)$

ALGORITHM 3. Stochastic query-by-committee algorithm. Data to query are chosen by sampling according to a probability distribution that gives more weights to data points whose predictions are uncertain. Uncertainty is inferred by a standard query-by-committee algorithm.

Fix the number of learners n in the ensemble.
Result: B elements from P^U
Input: $P^U, S^{(t)}, B$
 a) Train an ensemble $\{T_i\}_{i=1}^n$ of n models on data $S^{(t)}$.
 b) Compute predictions $\hat{E}_i = T_i(X) \forall X \in P^U, \forall i$
 c) Compute the community disagreement
 $q(X) = \text{std}(\hat{E}_i) \forall X \in P^U$
 d) Compute the weights: $L(X) = \frac{q(X) - q_{\min}}{q_{\max} - q_{\min}}$, and the sampling probability $p(X) = \frac{L(X)}{\sum_X L(X)}$ where $q_{\min} = \min_X q(X)$ and $q_{\max} = \max_X q(X)$
 e) sample B elements from the unlabeled data with probabilities $p(X)$

query,⁶⁸ which is very similar in spirit to Algorithm 3. However, we empirically observed Algorithm 3 to work better than this approach.

The balance between sampling points from the sparse and high-density regions is controlled by the function L , which is linear with respect to the community disagreement. The probability of a point being sampled decreases linearly with the decrease in the point's uncertainty. One can have more freedom on this balance by considering powers of this function, i.e., L^α , where $\alpha \in \mathbb{R}^+$. For $\alpha \in (0, 1)$, the algorithm will sample more points with low uncertainty and conversely less for $\alpha \in (1, \infty)$. We performed a heuristic study of the effect of different powers α . At each AL iteration, we ran SQBF algorithm for different values of $\alpha \in \{0.5, 0.75, 1, 1.25, 1.75\}$. We picked the α that led to the largest improvement in generalization error and collected the corresponding queried points. We proceeded to active learning using this batch as a part of the pool. The whole procedure was repeated at every AL iteration. We found only minor improvements of the accuracy when using multiple, optimized, values of α . We explored a few other heuristics of similar nature, but none of them yielded significantly better results. Hence, throughout the paper, we report results obtained with a single value of $\alpha = 1$. In Algorithms 2 and 3, it remains to specify the ensemble of learners and to elaborate on how to diversify them. While any ensemble of ML models can be used, we propose to use the trees of an RFR as members of this ensemble. In Sec. II B, we argue that choosing regression trees for inferring uncertainty is advantageous because of relatively low-training complexity and a straightforward diversification-ability. The RFR combined with Algorithm 3 gives rise to a regression version of the stochastic query by forest algorithm (SQBF),⁶⁶ employed in this study.

B. Random forest regressor

Regression trees are a non-parametric way of solving a regression problem. They are based on the intuition that the output value can be inferred by partitioning the input space. In particular, for solving a regression problem with data pairs $\{(X_i, E_i)\}_{i=1}^I$, a *tree regressor* aims at finding J distinct and non-overlapping regions

$R_1 \dots R_J$ in the feature space that minimize

$$\sum_{j=1}^J \sum_{X_i \in R_j} (E_i - \bar{E}_{R_j})^2,$$

where $\bar{E}_{R_j} = \sum_{X_i \in R_j} E_i / m_j$ is the average energy of m_j molecular geometries in region R_j .⁶⁹ This problem is nondeterministic polynomial-time- (NP-) complete.⁷⁰ Therefore, only near-optimal solutions are considered by restricting ourselves to hyper-rectangular regions and using recursive binary splitting, a greedy algorithm to obtain a near-optimal segmentation. A prediction for a new molecular geometry is done by assigning the input geometry to one of the regions. The prediction for this geometry is then the average energy of all geometries in the training dataset that fall in this region. A major drawback of tree regressors is their large variance.⁶⁹ A powerful approach to mitigate this problem is to consider an ensemble of trees. The key idea is that averaging a set of independent random variables, which have comparable variances, reduces their overall variance.⁶⁹ In an ensemble method, a random perturbation is introduced to the learning process in order to produce several different learners from the same training set. Thus, taking the average of the predictions of the ensemble would result in a reduction in variance. Such a random perturbation can be introduced by bootstrapping, which gives rise to bootstrap aggregation (bagging) methods. Here, B different bootstrapped datasets of size m_b are generated. A tree is built on each model. For a new data point, a prediction is made by taking the average of the predictions of all trees: $T_{\text{bag}}(x) = \frac{1}{B} \sum_{i=1}^B T_i(x)$.

A further random perturbation in tree models can be introduced by considering, at each split, only a randomly drawn subset of all possible features. This gives rise to the random forest regressor (RFR).⁷¹ Thus, we see that RFR employs a twofold randomization procedure. The ensemble can be made even more diverse by introducing further randomization in the learning process, e.g., extremely randomized trees.⁷² RFR is an inherent ensemble method encompassing diverse learners. This makes the model a very attractive option for a QBC-based algorithm such as Algorithms 2 and 3.

Another advantage of using the trees of an RFR in Algorithms 2 and 3 is its relatively low-training complexity. An AL paradigm is a dynamic paradigm that needs to be performed iteratively until we are satisfied with the performance. One wishes to be able to perform these iterations quickly. Otherwise, the time saved from performing redundant electronic structure calculations would be wasted in performing AL iterations. Building an RFR is relatively cheap. It has an average time complexity of $\Theta(M \cdot K \cdot \tilde{N} \log_2 \tilde{N})$,⁷³ where K and M denote the number of random features sampled at each splitting and the number of trees, respectively, and $\tilde{N} \approx 0.632N$ with the number N of training examples.⁷⁴ This should be compared to the computational cost of training a GP, which scales as $\mathcal{O}(N^3)$.⁴⁰ With extremely randomized trees, the average time complexity for training is $\Theta(M \cdot K \cdot N \log_2 N)$.⁷² The average inference complexity of RFR is $\Theta(M \cdot \log N)$.⁷² Thus, one AL iteration scales as $\mathcal{O}(M \cdot K \cdot \tilde{N} \log_2 \tilde{N})$ with the number of so-far-labeled data N . The complexity of RFR is asymptotically inferior to that of a NN, which has a training time complexity⁷⁵ of $\mathcal{O}(N_e \cdot N(\sum_{i=1}^{l-1} N^i \cdot N^{i+1}))$ with the number of epochs N_e needed for the NN to converge and the number of neurons N^i in layer i . However, in the data size regime of our

application, the computational costs of an RFR are smaller than that of the NN.

III. SIMULATIONS

We applied different policy algorithms and tested their performance for building the prototypical intermolecular PES of the pyrrole(H_2O) complex. Due to the highly fluxional nature of the hydrogen bond, the intermolecular motions in pyrrole(H_2O) are highly delocalized, rendering the calculation and representation of the PES very challenging. The intramolecular vibrations in the pyrrole and water moieties can be described with a relatively simple, though multi-dimensional, single-minimum PES and thus, for simplicity of calculations, were not considered here. We fixed the structures of pyrrole and water monomers to the experimentally determined values^{76,77} (see the [supplementary material](#)) and varied the six intermolecular coordinates, as shown in Fig. 1. These are defined as follows: the relative position of water with respect to pyrrole is described by the three spherical coordinates $R = [0.2, 1]$ nm, $\theta = [0, \pi]$, and $\phi = [0, \pi]$ and the relative orientation of water is defined by the three Euler angles $\alpha = [0, \pi]$, $\beta = [0, \pi]$, and $\gamma = [0, \pi]$. The angles ϕ , α , and γ were restricted to the ranges $[0, \pi]$ exploiting the $C_{2v}(\text{M})$ symmetry of the complex.

The pool of molecular configurations was generated *a priori* as the direct product of one-dimensional grids for every degree of freedom and contained 57 500 different molecular geometries covering the potential energy up to 5000 cm^{-1} above dissociation. All coordinates were sampled more densely in the vicinity of the equilibrium geometry. In addition, the angular coordinates were sampled more densely for small radial distances $R \leq 500$ pm with a sparser grid for $500 < R \leq 1000$ pm. This led to a nonuniform distribution of energies in the pool, as shown in Fig. 2. Note that a direct-product grid is not essential for the accumulation of the pool of unlabeled geometries and the test dataset. Here, it was used mainly because it allows the coverage of the whole configuration space that is relevant for the subsequent quantum dynamics' simulations, and hence prevents biases and holes in the pool and test data. While this method is not arbitrarily extendable to systems with more degrees of freedom, other pool accumulation methods⁵¹ could be used without modifications to the SQBF approach. The electronic structure calculations employed the density-fitting explicitly correlated




FIG. 1. Internal intermolecular coordinates R , θ , ϕ , α , β , and γ of pyrrole(H_2O).




FIG. 2. The probability density distribution of the energies corresponding to all molecular geometries in the pool. The histogram was calculated for a bin width of 34.5 cm^{-1} and has a peak at 1600 cm^{-1} , corresponding to the dissociation limit of pyrrole(H_2O).

DF-MP2-F12 level of theory^{78–80} in the frozen-core approximation using aug-cc-pVDZ-F12⁸¹ atomic orbital, cc-pVDZ-F12+/OPTRI⁸² resolution of the identity, and aug-cc-pVDZ/MP2FIT⁸³ density fitting basis sets. The geminal exponent was fixed at 1.0. The electronic structure calculations were carried out using Molpro.^{84–86} A subset of 10% of the total number of points in the pool was randomly selected as a test set and taken out of the pool (OOP). 5% of the remaining data was randomly selected as a validation set. We employed two different machine learning models, RFR and NN, to fit the data. We used exponential functions of interatomic distances, with all distances considered, as molecular descriptors; see the [Appendix](#).

Furthermore, based on reviewer suggestions, we tested the SQBF algorithm and an NN model on the PES of the N_4 molecule using previously reported electronic structure data.⁸⁷

A. Performance

For pyrrole(H_2O), we compared the performance of the RS, QBC, and SQBF AL policy algorithms considering the convergence rate and the fitting accuracy. For QBC (Algorithm 2) and SQBF (Algorithm 3), we used the trees of an RFR as an ensemble of learners. All policy algorithms started from the same fixed amount of $m = 2458$ labeled samples and queried the same equal number of m samples at every AL iteration. For every iteration and query algorithm, we used the RFR and NN models to fit the data. The fitting error is defined as the rms error of the ML models in predicting the energies on the OOP dataset. This dataset was the same for all policy algorithms and followed the joint distribution of the problem \mathcal{P} . The accuracy of a model on this dataset is an estimate of the generalization error. Further technical details on the training process of the RFR and NN models and their hyper parameters are provided in the [Appendix](#).

The fitting errors of the RFR and NN models for different policy algorithms are plotted in [Fig. 3](#) as functions of AL iteration. The SQBF strategy with the RFR model leads to the fastest convergence of the error. The QBC strategy outperforms RS. Similar convergence behavior of different policy algorithms can be observed for the NN model. For our dataset, the fitting error of NN was smaller than that of RFR for all AL iterations and for all strategies by an average factor of 3.3. [Table I](#) summarizes these results. The better performance of NNs is partially due to the fact that NNs are easier to train to higher




FIG. 3. rms error of out-of-the-pool datasets using (up) the random forest regressor and (down) the neural network models for the RS (triangles), QBC (circles), and SQBF (squares) policies. The SQBF has the fastest convergence. A neural network model, trained on 30% of the total amount of data points in the pool, achieves an rms error of 16 cm^{-1} . The rms error on the full dataset is 11 cm^{-1} . The neural networks trained on data collected by the QBC or RS algorithms show worse performance. The same convergence patterns hold when using a random forest regressor to train on the data instead of a neural network, albeit at overall somewhat slower convergence.

accuracy and can approximate complex functions with a better control on the bias-variance trade-off, which was enabled by using an early stopping criterion on the validation set; see the [Appendix](#). The AL iterations were terminated when the pool became empty. In practice, the iterations are to be terminated when the derivative of the fitting error with respect to the amount of labeled data is less than a predefined value⁵⁴ or simply when the fitting error of the model is small enough.

Similarly, for the N_4 molecule, the SQBF algorithm was used to query geometries from the pool of 16 421 molecular geometries reported.⁸⁷ The OOP and validation datasets were each generated using 10% of the uniform-randomly sampled pool data. An initial batch of 240 geometries was uniform-randomly sampled from the pool and the SQBF algorithm queried 240 geometries at each AL iteration. The same molecular descriptor as described above was used to transform the data and an NN model was used for fitting;

TABLE I. Out-of-the-pool rms errors (in cm^{-1}) of the random forest regressor and neural network models, listed as RFR/NN, computed for various fractions of the total pool data collected by the different AL policies.

AL policy	20%	40%	60%	80%	100%
RS	183/57	117/31	81/20	52/15	39/11
QBC	141/43	74/21	49/13	41/11	38/11
SQBF	88/27	37/14	36/12	38/11	39/11

TABLE II. rms mean errors and standard deviations using the available data (in cm^{-1}) of NN using the AL SQBF algorithm with a NN to fit the data (present work) and latin hypercube sampling with GPs to fit the data,⁸⁸ applied to PES data of the N_4 molecule.⁸⁷

No. training data	NN (cm^{-1})	GP (cm^{-1})
240	$36\,518 \pm 2\,697$	$13\,300 \pm 2\,770$
480	$26\,207 \pm 1\,871$	$10\,027 \pm 1\,371$
720	$11\,192 \pm 1\,200$	$8\,401 \pm 1\,102$
960	$8\,111 \pm 668$	$7\,544 \pm 972$
1200	$6\,201 \pm 462$	$6\,806 \pm 962$
1680	$4\,704 \pm 612$...
1800	$4\,494 \pm 633$	$5\,551 \pm 951$
1920	$4\,284 \pm 658$...
2400	$3\,557 \pm 675$	$5\,012 \pm 832$

details on the NN design are provided in Sec. I. This procedure was repeated 100 times and the mean and standard deviation of the resulting NN errors on the entire dataset as a function of the number of training examples are reported in Table II. The SQBF

results are compared with the ensemble of 100 GPs used to fit the data collected by the Latin hypercube sampling algorithm.⁸⁸ The GP method shows a better performance for the first few AL iterations. We attribute this to the fact that it is hard to prevent overfitting with a neural network with a very small set of randomly selected training data. However, already at 1200 training points, the two models result in comparable accuracy. With 1680 training points, our SQBF/NN approach achieves the same accuracy as GP with 2400 points, which corresponds to a 30% reduction in the size of the training dataset. All the following further investigations are performed for pyrrole(H_2O).

B. Distribution of queried data

In Fig. 4, we plotted the normalized distributions of the samples' electronic energies of pyrrole(H_2O) collected by different AL policies at three different iterations corresponding to 20%, 40%, and 60% of the total pool. We compare these with the distribution of energies in the total pool in Fig. 2, which has a peak around 1600 cm^{-1} , corresponding to the dissociation limit of pyrrole(H_2O). The densities were computed using 200 equally sized bins covering

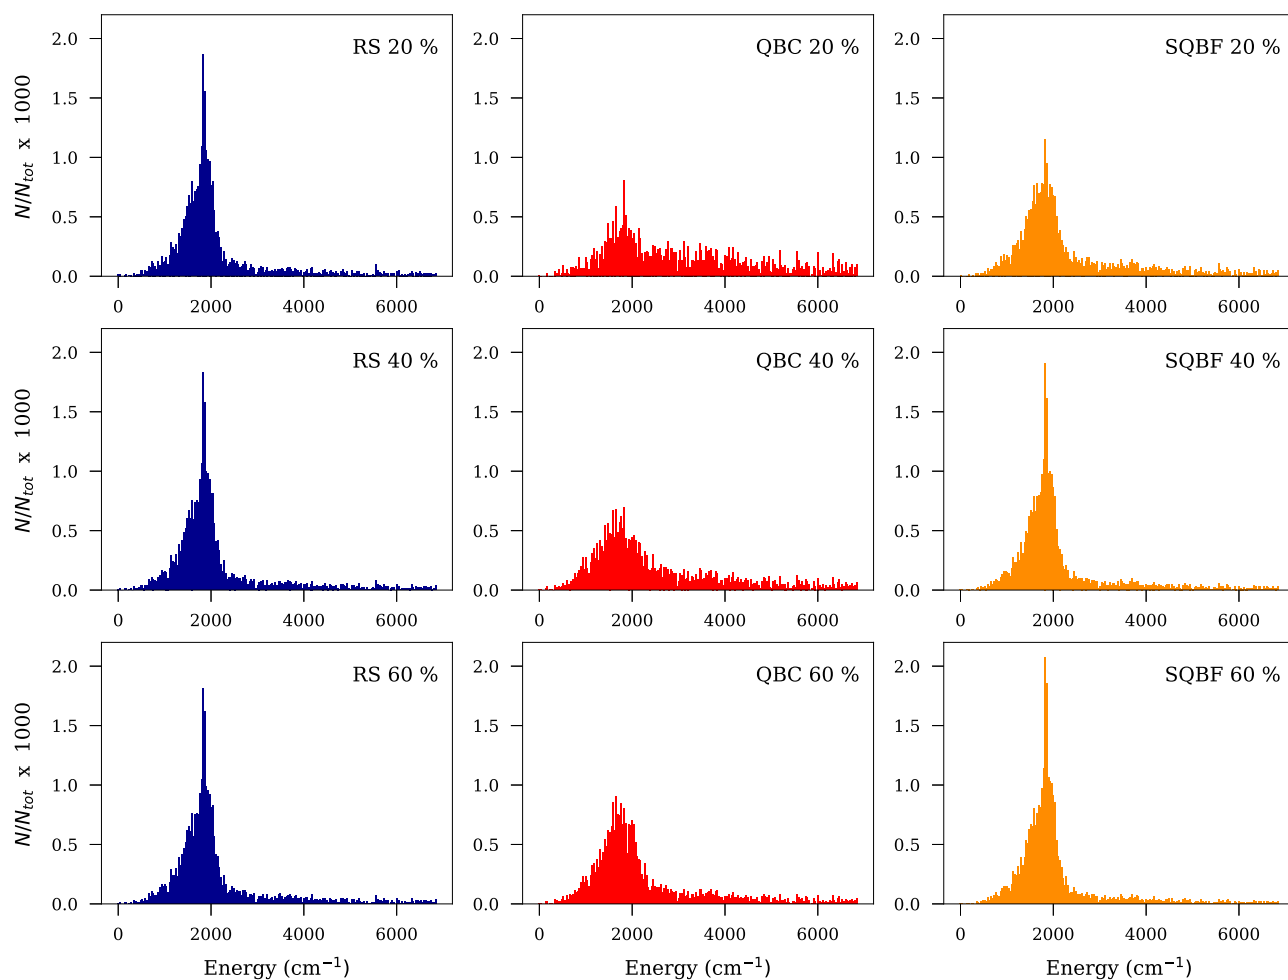


FIG. 4. Normalized probability density distributions of the number of data points N/N_{tot} across the potential energies plotted for the data collected by the RS, QBC, and SQBF policies at different AL iterations corresponding to 20%, 40%, and 60% of the total pool. The bin width of the histograms is 34.5 cm^{-1} .

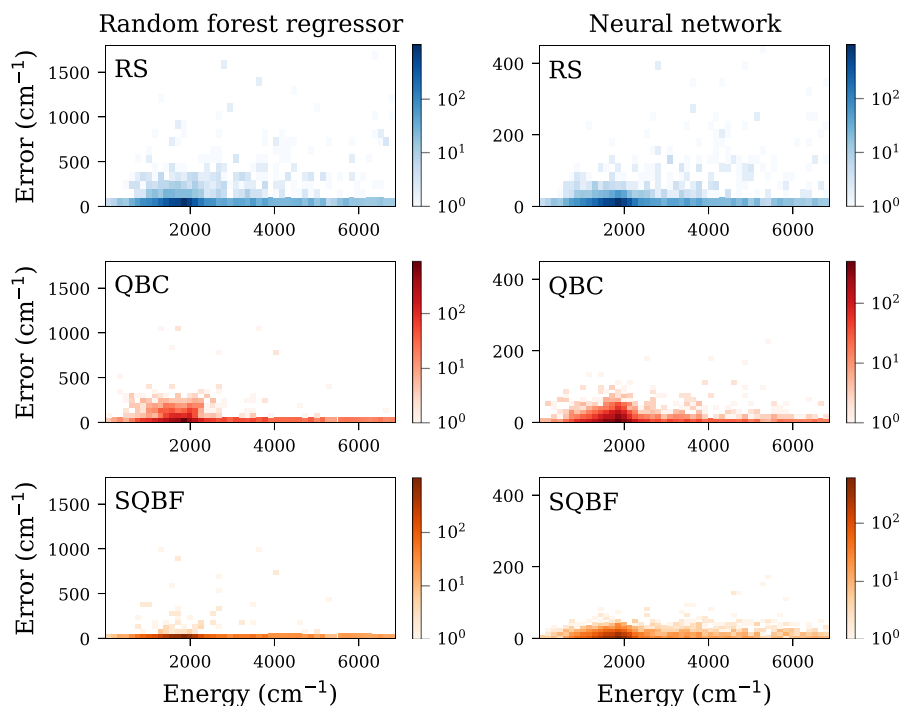


FIG. 5. 2D histograms of discrepancies between the predictions of random forest regressor and neural network models (trained on 40% of the pool) and the potential energy of the out-of-the-pool data for different policy algorithms; 20 and 50 bins were used for energy and absolute error, respectively. Models trained on data collected by QBC tend to perform better on high-energy regions than on low-energy regions. The opposite is true for RS. In contrast, models trained on data collected by SQBF have a more uniform accuracy across the whole energy spectrum.

the energy range from 0 to 6874 cm^{-1} and normalized to the bin width of 34.5 cm^{-1} . Evidently, the probability density of data sampled by the RS policy most closely resembles the pool distribution. On the other hand, it is clear that the QBC algorithm samples more data with higher energies, whereas SQBF keeps a balance between both the RS and QBC tendencies. As the number of the labeled data increases, all probability density distributions become more similar to the distribution in the pool.

It is reasonable to expect that a model built on a dataset sampled by the QBC algorithm will tend to have a better performance for the high-energy regions. This is demonstrated in Fig. 5 showing the 2D histograms of OOP energies and the absolute errors of the RFR and NN models in predicting these energies, plotted for different policy algorithms. The histograms were computed using 20 and 50 equally sized bins for the energy and absolute errors, respectively. The size of the training dataset here corresponds to 40% of the pool's size. We clearly see that RS achieves good accuracy for the points with low energies, QBC works best for the points with high energies, and the SQBF maintains a more regular accuracy across the whole energy spectrum.

C. Dependency on batch size and size of initially labeled dataset

We repeated the above calculations with a smaller batch size of 122 points instead of initially used 2458, starting from the same initially labeled dataset. The convergence of the RFR fitting error with the number of training data is plotted in Fig. 6 for different policy algorithms. Here, we note that both QBC and SQBF strategies benefit slightly from using a smaller batch size. This is in accordance with previous studies that showed a decreasing performance of QBC

with increasing batch size, which is due to collecting many similar samples.⁸⁹

We also studied the effect of changing the size of initially labeled dataset. Figure 7 shows the RFR fitting errors for different policy algorithms obtained from initial datasets of 100 and 2458 samples with the batch size of 2458. We observe that the RS policy algorithm outperforms QBC and that the accuracy of QBC declines significantly. This suggests that with a fewer number of initially labeled data, an AL strategy should focus on collecting grid points from dense regions of the configuration space rather than sampling points with high uncertainties in their predictions. Notably, the SQBF performance is not affected by the size change.




FIG. 6. Effect of the batch size on the out-of-the-pool error of an RFR model trained using data collected by the RS (blue, triangles), QBC (red, circles), and SQBF (orange, squares) policies. Solid (points) and dashed lines correspond to the batch sizes fixed to 2458 and 122, respectively.




FIG. 7. Effect of the size of initially labeled data on the out-of-the-pool error of an RFR model trained using data collected by the RS (blue triangles), QBC (red circles), and SQBF (orange squares) policies. Solid (points) and dashed lines correspond to 100 and 2458 initially labeled data, respectively.

IV. SUMMARY AND CONCLUSION

The first principles calculations of molecular PESs, especially for molecules with many fluxional degrees of freedom, are computationally expensive. One of the major bottlenecks originates from the need to perform the expensive quantum chemical calculations for tens and hundreds of thousands of different molecular geometries. Algorithms that allow one to reduce the number of necessary single-point calculations with controlled accuracy of the resulting PES are thus highly demanded. For small molecules, grid reduction algorithms were found beneficial in calculations employing high-level electron correlation, basis set, and relativistic corrections, which are usually computationally affordable only for a relatively small number of points.^{90–93}

We developed an accurate and automated procedure to efficiently sample the molecular geometry grid points, leading to a systematic convergence of the accuracy of PES. We employed a regression version of SQBF, a pool-based active learning algorithm to generate a compact grid of molecular geometries and the RFR and NN ML models to construct the six-dimensional intermolecular PES of the weakly bound pyrrole(H_2O) complex. The proposed method samples grid points with high uncertainties in the corresponding predictions of energies and, at the same time, preserves the statistical information embedded in the pool. In our benchmark application to the six-dimensional intermolecular PES of pyrrole(H_2O), this led to a roughly two times faster convergence with respect to the number of grid points than the commonly used QBC algorithm to represent the PES to an accuracy of about 16 cm^{-1} . Furthermore, the PES fitted on the data sampled by SQBF exhibited a more uniform accuracy across the whole energy spectrum in comparison to QBC. We empirically showed that the SQBF method is not very sensitive to a variation of parameters such as the size of initially labeled data and size of the batch.

In addition, the proposed method is computationally cheap and scales well with the size of the labeled data N , i.e., as $\mathcal{O}(M \cdot K \cdot \tilde{N} \log^2 \tilde{N})$, where K and M denote the number of random features sampled at each splitting and the number of trees, respectively, $\tilde{N} = 0.632N$. This makes the method attractive for developing universal ML-potentials where large datasets are needed.^{94–96} In the

case when the accuracy of the RFR is not sufficient for the application of the PES, we showed that the data can be used equally well by other ML models such as NNs. An alternative would be to employ Algorithm 3 with any other ensemble of learners or even with a model that offers a direct computation of uncertainty.

Overall, the presented procedure is general and can be applied to the PESs of any polyatomic molecule. It can also be used to model other physical properties such as dipole-moment or polarizability surfaces. The major advantage of the proposed method over the more popular QBC approach is the heuristic sampling procedure that preserves the distribution of data in the pool while keeping the uncertainty as the primal selection criterion. We believe that in the future, the general approach can be improved even further by a better tuned balance between uncertainty and representativeness.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for the computed potential energies of pyrrole(H_2O) and the structures of pyrrole and water monomers.

ACKNOWLEDGMENTS

This work was supported by Deutsches Elektronen-Synchrotron DESY, a member of the Helmholtz Association (HGF), the Data Science in Hamburg HELMHOLTZ Graduate School for the Structure of Matter (DASHH, No. HIDSS-0002), and by the Deutsche Forschungsgemeinschaft (DFG) through the cluster of excellence “Advanced Imaging of Matter” (AIM, EXC 2056, ID 390715994). The authors acknowledge the use of the Maxwell computational resources operated at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available within the article and its [supplementary material](#). The computer code and all relevant data for this work can be obtained from the git repository at <https://github.com/CFEL-CMI/Active-Learning-of-PES>.

APPENDIX: TECHNICAL DETAILS ON THE RANDOM PERTURBATIONS IN AL ALGORITHMS AND THE TRAINING OF MACHINE LEARNING MODELS

The regression trees used to implement the QBC and SQBF algorithms were both built using the scikit-learn (`sklearn`) Python package.⁹⁷ All AL algorithms used here were written based on the Libact Python package.⁹⁸ For both the QBC and SQBF algorithms, we used an ensemble of 100 trees. The training during all AL iterations used an exponential function of the inter-nuclear distances as a molecular descriptor, $1 - \exp(-(r - r_0))$, where r and r_0 denote the

actual distance and equilibrium distance between two atoms, respectively. The perturbation of the learning process is controlled through two parameters: (i) a bootstrapping parameter γ that determines the fraction of data sampled by each tree and (ii) the number of features β sampled randomly by each tree. For the batch sizes used in the simulations, we experimented with several combinations of these parameters and obtained the best convergence for $\gamma = 0$ and $\beta = 12$ for simulations on pyrrole(H_2O) and $\gamma = 0$ and $\beta = 4$ for simulations on N_4 . The same parameters were used for uncertainty estimation in both the QBC and SQBF algorithms. Minimal cost complexity pruning was used to reduce the overfitting of RFR with complexity parameter $c = 0.01$. Since we require that Algorithm 3 queries exactly B geometries, one may run into a situation where the number of entries with non-zero probabilities of the distribution p is less than B . In such a case, we chose to query the elements with highest uncertainty. The effect of this choice on the simulations conducted in the manuscript is negligible since we only encountered this case once in one of the last AL iterations.

The NN used is a multilayer perceptron, and training was implemented using the Python Tensorflow package.⁹⁹ The NN has three hidden layers with 256, 512, and 256 neurons, respectively, and a single neuron output layer. The second and third layers were l_2 -regularized with a regularization parameter of 10^{-5} . All hidden layers used "ReLU" as the activation function. The ReLU activation function could be substituted with a smooth approximation such as Softplus to yield a smooth PES. The same aforementioned molecular descriptor was used. The networks were trained for 250 epochs using the Adam optimization algorithm,¹⁰⁰ with an initial learning rate of 0.0025 and a decaying learning rate schedule ($\text{lr}_{\text{current}} = 0.9825 \times \text{lr}_{\text{previous}}$). An early stopping callback was employed on the validation set that was taken out of the pool with a patience of 25. The NN hyper parameters were set to obtain a sufficiently accurate NN, with a test error of around 10 cm^{-1} when using all the training data.

REFERENCES

- K. Liu, "Quantum dynamical resonances in chemical reactions: From A + BC to polyatomic systems," *Adv. Chem. Phys.* **149**, 1–46 (2012).
- A. G. Suits, "Roaming reactions and dynamics in the van der Waals region," *Annu. Rev. Phys. Chem.* **71**, 77–100 (2020).
- M. Colombo, D. Rau, and V. Parsegian, "Protein solvation in allosteric regulation: A water effect on hemoglobin," *Science* **256**, 655–659 (1992).
- T. S. Zwier, "The spectroscopy of solvation in hydrogen-bonded aromatic clusters," *Annu. Rev. Phys. Chem.* **47**, 205–241 (1996).
- T. M. Korter, D. W. Pratt, and J. Küpper, "Indole- H_2O in the gas phase. Structures, barriers to internal motion, and $S_1 \leftarrow S_0$ transition moment orientation. Solvent reorganization in the electronically excited state," *J. Phys. Chem. A* **102**, 7211–7216 (1998).
- I. V. Hertel and W. Radloff, "Ultrafast dynamics in isolated molecules and molecular clusters," *Rep. Prog. Phys.* **69**, 1897–2003 (2006).
- C. Kang and D. W. Pratt, "Structures, charge distributions, and dynamical properties of weakly bound complexes of aromatic molecules in their ground and electronically excited states," *Int. Rev. Phys. Chem.* **24**, 1–36 (2007).
- A. K. Samanta, Y. Wang, J. S. Mancini, J. M. Bowman, and H. Reisler, "Energetics and predissociation dynamics of small water, HCl, and mixed HCl-water clusters," *Chem. Rev.* **116**, 4913–4936 (2016).
- M. Johny, C. A. Schouder, A. Al-Refaie, L. He, J. Wiese, H. Stapelfeldt, S. Trippel, and J. Küpper, "Molecular sunscreen: Water protects pyrrole from radiation damage" (submitted); [arXiv:2010.00453](https://arxiv.org/abs/2010.00453) [physics].
- J. Onvlee, S. Trippel, and J. Küpper, "Ultrafast light-induced dynamics in solvated biomolecules: The indole chromophore with water" (unpublished); [arXiv:2103.07171](https://arxiv.org/abs/2103.07171) [physics].
- O. Akin-Ojo and K. Szalewicz, "Potential energy surface and second virial coefficient of methane-water from *ab initio* calculations," *J. Chem. Phys.* **123**, 134311 (2005).
- M. P. Metz, K. Szalewicz, J. Sarka, R. Tóbiás, A. G. Császár, and E. Mátyus, "Molecular dimers of methane clathrates: *Ab initio* potential energy surfaces and variational vibrational states," *Phys. Chem. Chem. Phys.* **21**, 13504–13525 (2019).
- Q. Ma and H.-J. Werner, "Accurate intermolecular interaction energies using explicitly correlated local coupled cluster methods [PNO-LCCSD(T)-F12]," *J. Chem. Theory Comput.* **15**, 1044–1052 (2019).
- M. P. Metz and K. Szalewicz, "Automatic generation of flexible-monomer intermolecular potential energy surfaces," *J. Chem. Theory Comput.* **16**, 2317–2339 (2020).
- S. N. Vogels, T. Karman, J. Klos, M. Besemer, J. Onvlee, A. van der Avoird, G. C. Groenenboom, and S. Y. T. van de Meerakker, "Scattering resonances in bimolecular collisions between NO radicals and H_2 challenge the theoretical gold standard," *Nat. Chem.* **10**, 435–440 (2018).
- B. Brauer, M. K. Kesharwani, S. Kozuch, and J. M. L. Martin, "The $S66 \times 8$ benchmark for noncovalent interactions revisited: Explicitly correlated *ab initio* methods and density functional theory," *Phys. Chem. Chem. Phys.* **18**, 20905–20925 (2016).
- J. Behler, "Perspective: Machine learning potentials for atomistic simulations," *J. Chem. Phys.* **145**, 170901 (2016).
- K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, "Assessment and validation of machine learning methods for predicting molecular atomization energies," *J. Chem. Theory Comput.* **9**, 3404–3419 (2013).
- S. Manzhos and T. Carrington, "Neural network potential energy surfaces for small molecules and reactions," *Chem. Rev.* **121**, 10187 (2021).
- A. Blum, J. Hopcroft, and R. Kannan, *Foundations of Data Science* (Cambridge University Press, 2020).
- M. Lotz, "Mathematics of machine learning," 2018, lecture notes; accessed 30 March 2021.
- Y. Yao, "A mathematical introduction to data science," 2019, lecture notes; accessed 30 March 2021.
- S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, 2014).
- I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep Learning* (MIT Press, 2016), Vol. 1.
- X. Wang, P. L. Houston, and J. M. Bowman, "A new (multi-reference configuration interaction) potential energy surface for H_2CO and preliminary studies of roaming," *Philos. Trans. R. Soc., A* **375**, 20160194 (2017).
- B. J. Braams and J. M. Bowman, "Permutationally invariant potential energy surfaces in high dimensionality," *Int. Rev. Phys. Chem.* **28**, 577–606 (2009).
- Z. Xie and J. M. Bowman, "Permutationally invariant polynomial basis for molecular energy surface fitting via monomial symmetrization," *J. Chem. Theory Comput.* **6**, 26–34 (2010).
- C. Qu, Q. Yu, and J. M. Bowman, "Permutationally invariant potential energy surfaces," *Annu. Rev. Phys. Chem.* **69**, 151–175 (2018).
- R. Conte, C. Qu, P. L. Houston, and J. M. Bowman, "Efficient generation of permutationally invariant potential energy surfaces for large molecules," *J. Chem. Theory Comput.* **16**, 3264–3272 (2020).
- T. Morawietz, V. Sharma, and J. Behler, "A neural network potential-energy surface for the water dimer based on environment-dependent atomic energies and charges," *J. Chem. Phys.* **136**, 064103 (2012).
- J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," *Phys. Rev. Lett.* **98**, 146401 (2007).
- O. T. Unke and M. Meuwly, "PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges," *J. Chem. Theory Comput.* **15**, 3678–3693 (2019).
- S. Manzhos, R. Dawes, and T. Carrington, "Neural network-based approaches for building high dimensional and quantum dynamics-friendly potential energy surfaces," *Int. J. Quantum Chem.* **115**, 1012–1020 (2014).

- ³⁴J. Behler, "Constructing high-dimensional neural network potentials: A tutorial review," *Int. J. Quantum Chem.* **115**, 1032–1050 (2015).
- ³⁵B. Jiang, J. Li, and H. Guo, "Potential energy surfaces from high fidelity fitting of *ab initio* points: The permutation invariant polynomial - neural network approach," *Int. Rev. Phys. Chem.* **35**, 479–506 (2016).
- ³⁶C. Schran, J. Behler, and D. Marx, "Automated fitting of neural network potentials at coupled cluster accuracy: Protonated water clusters as testing ground," *J. Chem. Theory Comput.* **16**, 88–99 (2019).
- ³⁷A. Kamath, R. A. Vargas-Hernández, R. V. Krems, T. Carrington, and S. Manzhos, "Neural networks vs Gaussian process regression for representing potential energy surfaces: A comparative study of fit quality and vibrational spectrum accuracy," *J. Chem. Phys.* **148**, 241702 (2018).
- ³⁸A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, "Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons," *Phys. Rev. Lett.* **104**, 136403 (2010).
- ³⁹C. Qu, Q. Yu, B. L. Van Hoozen, J. M. Bowman, and R. A. Vargas-Hernández, "Assessing Gaussian process regression and permutationally invariant polynomial approaches to represent high-dimensional potential energy surfaces," *J. Chem. Theory Comput.* **14**, 3381–3396 (2018).
- ⁴⁰C. E. Rasmussen and C. K. I. Williams, in *Gaussian Processes for Machine Learning*, edited by T. Dietterich (The MIT Press, Cambridge, MS, 2006).
- ⁴¹H. Sugisawa, T. Ida, and R. V. Krems, "Gaussian process model of 51-dimensional potential energy surface for protonated imidazole dimer," *J. Chem. Phys.* **153**, 114101 (2020).
- ⁴²J. Dai and R. V. Krems, "Interpolation and extrapolation of global potential energy surfaces for polyatomic systems by Gaussian processes with composite kernels," *J. Chem. Theory Comput.* **16**, 1386–1395 (2020).
- ⁴³R. A. Vargas-Hernández, Y. Guan, D. H. Zhang, and R. V. Krems, "Bayesian optimization for the inverse scattering problem in quantum reaction dynamics," *New J. Phys.* **21**, 022001 (2019).
- ⁴⁴O. T. Unke and M. Meuwly, "Toolkit for the construction of reproducing kernel-based representations of data: Application to multidimensional potential energy surfaces," *J. Chem. Inf. Model.* **57**, 1923–1931 (2017).
- ⁴⁵P. O. Dral, A. Owens, S. N. Yurchenko, and W. Thiel, "Structure-based sampling and self-correcting machine learning for accurate calculations of potential energy surfaces and vibrational levels," *J. Chem. Phys.* **146**, 244108 (2017).
- ⁴⁶D. Koner and M. Meuwly, "Permutationally invariant, reproducing kernel-based potential energy surfaces for polyatomic molecules: From formaldehyde to acetone," *J. Chem. Theory Comput.* **16**, 5474–5484 (2020).
- ⁴⁷A. Iske, *Approximation Theory and Algorithms for Data Analysis* (Springer International Publishing, 2018).
- ⁴⁸B. Settles, "Active learning literature survey," Technical Report No. 1648, Department of Computer Sciences, University of Wisconsin-Madison, 2009.
- ⁴⁹A. A. Peterson, R. Christensen, and A. Khorshidi, "Addressing uncertainty in atomistic machine learning," *Phys. Chem. Chem. Phys.* **19**, 10978–10985 (2017).
- ⁵⁰Q. Lin, Y. Zhang, B. Zhao, and B. Jiang, "Automatically growing global reactive neural network potential energy surfaces: A trajectory-free active learning strategy," *J. Chem. Phys.* **152**, 154104 (2020).
- ⁵¹L. Zhang, D.-Y. Lin, H. Wang, R. Car, and W. E, "Active learning of uniformly accurate interatomic potentials for materials simulation," *Phys. Rev. Mater.* **3**, 023804 (2019).
- ⁵²E. Uteva, R. S. Graham, R. D. Wilkinson, and R. J. Wheatley, "Active learning in Gaussian process interpolation of potential energy surfaces," *J. Chem. Phys.* **149**, 174114 (2018).
- ⁵³T. D. Loeffler, T. K. Patra, H. Chan, M. Cherukara, and S. K. R. S. Sankaranarayanan, "Active learning the potential energy landscape for water clusters from sparse training data," *J. Phys. Chem. C* **124**, 4907–4916 (2020).
- ⁵⁴Y. Zhai, A. Caruso, S. Gao, and F. Paesani, "Active learning of many-body configuration space: Application to the Cs⁺-water MB-nrg potential energy function as a case study," *J. Chem. Phys.* **152**, 144103 (2020).
- ⁵⁵J. Vandermause, S. B. Torrisi, S. Batzner, Y. Xie, L. Sun, A. M. Kolpak, and B. Kozinsky, "On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events," *npj Comput. Mater.* **6**, 20 (2020).
- ⁵⁶E. V. Podryabinkin and A. V. Shapeev, "Active learning of linearly parametrized interatomic potentials," *Comput. Mater. Sci.* **140**, 171–180 (2017).
- ⁵⁷M. Gastegger, J. Behler, and P. Marquetand, "Machine learning molecular dynamics for the simulation of infrared spectra," *Chem. Sci.* **8**, 6924–6935 (2017).
- ⁵⁸J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev, and A. E. Roitberg, "Less is more: Sampling chemical space with active learning," *J. Chem. Phys.* **148**, 241733 (2018).
- ⁵⁹G. Sivaraman, A. N. Krishnamoorthy, M. Baur, C. Holm, M. Stan, G. Csányi, C. Benmore, and Á. Vázquez-Mayagoitia, "Machine-learned interatomic potentials by active learning: Amorphous and liquid hafnium dioxide," *npj Comput. Mater.* **6**, 104 (2020).
- ⁶⁰Y. Guan, S. Yang, and D. H. Zhang, "Construction of reactive potential energy surfaces with Gaussian process regression: Active data selection," *Mol. Phys.* **116**, 823–834 (2018).
- ⁶¹Q. Lin, L. Zhang, Y. Zhang, and B. Jiang, "Searching configurations in uncertainty space: Active learning of high-dimensional neural network reactive potentials," *J. Chem. Theory Comput.* **17**, 2691–2701 (2021).
- ⁶²H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proceedings of the Fifth Annual Workshop on Computational Learning Theory* (ACM, 1992), pp. 287–294.
- ⁶³S. Kee, E. del Castillo, and G. Runger, "Query-by-committee improvement with diversity and density in batch active learning," *Inf. Sci.* **454–455**, 401–418 (2018).
- ⁶⁴Z. Wang and J. Ye, "Querying discriminative and representative samples for batch mode active learning," *ACM Trans. Knowl. Discovery Data* **9**, 1–23 (2015).
- ⁶⁵C. Shui, F. Zhou, C. Gagné, and B. Wang, "Deep active learning: Unified and principled method for query and training," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, edited by S. Chiappa and R. Calandra (PMLR, Online, 2020), Vol. 108, pp. 1308–1318.
- ⁶⁶A. Borisov, E. Tuv, and G. Runger, "Active batch learning with stochastic query-by-forest (SQBF)," in *Active Learning and Experimental Design Workshop in Conjunction with AISTATS 2010* (PMLR, 2011), pp. 59–69.
- ⁶⁷M. Johny, J. Onvlee, T. Kierspel, H. Bieker, S. Trippel, and J. Küpper, "Spatial separation of pyrrole and pyrrole-water clusters," *Chem. Phys. Lett.* **721**, 149–152 (2019); [arXiv:1901.05267](https://arxiv.org/abs/1901.05267) [physics].
- ⁶⁸Y. Freund, H. S. Seung, E. Shamir, and N. Tishby, "Information, prediction, and query by committee," in *Advances in Neural Information Processing Systems* (Morgan Kaufmann Publishers, 1993), pp. 483–490.
- ⁶⁹G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning* (Springer, 2013), Vol. 112.
- ⁷⁰H. Laurent and R. L. Rivest, "Constructing optimal binary decision trees is NP-complete," *Inf. Process. Lett.* **5**, 15–17 (1976).
- ⁷¹L. Breiman, "Random forests," *Mach. Learn.* **45**, 5–32 (2001).
- ⁷²P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Mach. Learn.* **63**, 3–42 (2006).
- ⁷³G. Louppe, "Understanding random forests," Ph.D. thesis, University of Liège, 2014; [arXiv:1407.7502](https://arxiv.org/abs/1407.7502) [stat.ML].
- ⁷⁴The probability of not selecting a point in n draws of n samples with replacement is $(1 - 1/n)^n$, which converges in the limit of $n \rightarrow \infty$ to e^{-1} . Hence, bootstrap samples draw, on average, $1 - 1/e \approx 63.2\%$ of unique samples.^{73,101}
- ⁷⁵This bound can be straightforwardly obtained by noting that matrix multiplications are the most expensive computations in the forward and backward passes of the NN training. We assume here that matrix multiplication scales as $\Theta(N^3)$.
- ⁷⁶M. J. Tubergen, A. M. Andrews, and R. L. Kuczkowski, "Microwave spectrum and structure of a hydrogen-bonded pyrrole-water complex," *J. Phys. Chem.* **97**, 7451–7457 (1993).
- ⁷⁷U. Nygaard, J. T. Nielsen, J. Kirchheiner, G. Maltesen, J. Rastrup-Andersen, and G. Sørensen, "Microwave spectra of isotopic pyrroles: Molecular structure, dipole moment and ¹⁴N quadrupole coupling constants of pyrrole," *J. Mol. Struct.* **3**, 491–506 (1969).
- ⁷⁸H.-J. Werner and F. R. Manby, "Explicitly correlated second-order perturbation theory using density fitting and local approximations," *J. Chem. Phys.* **124**, 054114 (2006).
- ⁷⁹F. R. Manby, H.-J. Werner, T. B. Adler, and A. J. May, "Explicitly correlated local second-order perturbation theory with a frozen geminal correlation factor," *J. Chem. Phys.* **124**, 094103 (2006).

- ⁸⁰H.-J. Werner, T. B. Adler, and F. R. Manby, "General orbital invariant MP2-F12 theory," *J. Chem. Phys.* **126**, 164102 (2007).
- ⁸¹K. A. Peterson, T. B. Adler, and H.-J. Werner, "Systematically convergent basis sets for explicitly correlated wavefunctions: The atoms H, He, B–Ne, and Al–Ar," *J. Chem. Phys.* **128**, 084102 (2008).
- ⁸²R. A. Shaw and J. G. Hill, "Approaching the Hartree-Fock limit through the complementary auxiliary basis set singles correction and auxiliary basis sets," *J. Chem. Theory Comput.* **13**, 1691–1698 (2017).
- ⁸³F. Weigend, A. Köhn, and C. Hättig, "Efficient use of the correlation consistent basis sets in resolution of the identity MP2 calculations," *J. Chem. Phys.* **116**, 3175–3183 (2002).
- ⁸⁴H.-J. Werner, P. J. Knowles, F. R. Manby, J. A. Black, K. Doll, A. Heßelmann, D. Kats, A. Köhn, T. Korona, D. A. Kreplin, Q. Ma, T. F. Miller, A. Mitrushchenkov, K. A. Peterson, I. Polyak, G. Rauhut, and M. Sibaev, "The Molpro quantum chemistry package," *J. Chem. Phys.* **152**, 144107 (2020).
- ⁸⁵H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, and M. Schütz, "Molpro: A general-purpose quantum chemistry program package," *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2**, 242–253 (2012).
- ⁸⁶H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, M. Schütz, P. Celani, W. Györffy, D. Kats, T. Korona, R. Lindh, A. Mitrushchenkov, G. Rauhut, K. R. Shamasundar, T. B. Adler, R. D. Amos, S. J. Bennie, A. Bernhardtsson, A. Berning, D. L. Cooper, M. J. O. Deegan, A. J. Dobbyn, F. Eckert, E. Goll, C. Hampel, A. Hesselmann, G. Hetzer, T. Hrenar, G. Jansen, C. Köppl, S. J. R. Lee, Y. Liu, A. W. Lloyd, Q. Ma, R. A. Mata, A. J. May, S. J. McNicholas, W. Meyer, T. F. Miller III, M. E. Mura, A. Nicklass, D. P. O'Neill, P. Palmieri, D. Peng, K. Pflüger, R. Pitzer, M. Reiher, T. Shiozaki, H. Stoll, A. J. Stone, R. Tarroni, T. Thorsteinsson, M. Wang, and M. Welborn, MOLPRO, version, a package of *ab initio* programs. See <https://www.molpro.net>.
- ⁸⁷Y. Paukku, K. R. Yang, Z. Varga, and D. G. Truhlar, "Global *ab initio* ground-state potential energy surface of N₄," *J. Chem. Phys.* **139**, 044309 (2013).
- ⁸⁸J. Cui and R. V. Krems, "Efficient non-parametric fitting of potential energy surfaces for polyatomic molecules with Gaussian processes," *J. Phys. B: At. Mol. Opt. Phys.* **49**, 224001 (2016).
- ⁸⁹J. T. Ash, C. Zhang, A. Krishnamurthy, J. Langford, and A. Agarwal, "Deep batch active learning by diverse, uncertain gradient lower bounds," in International Conference on Learning Representations (ICLR), 2020.
- ⁹⁰K. A. Peterson, D. Feller, and D. A. Dixon, "Chemical accuracy in *ab initio* thermochemistry and spectroscopy: Current strategies and future challenges," *Theor. Chem. Acc.* **131**, 1079 (2012).
- ⁹¹A. Owens, S. N. Yurchenko, A. Yachmenev, J. Tennyson, and W. Thiel, "Accurate *ab initio* vibrational energies of methyl chloride," *J. Chem. Phys.* **142**, 244306 (2015).
- ⁹²A. Yachmenev, S. N. Yurchenko, T. Ribeyre, and W. Thiel, "High-level *ab initio* potential energy surfaces and vibrational energies of H₂CS," *J. Chem. Phys.* **135**, 074302 (2011).
- ⁹³P. O. Dral, A. Owens, A. Dral, and G. Csányi, "Hierarchical machine learning of potential energy surfaces," *J. Chem. Phys.* **152**, 204110 (2020).
- ⁹⁴R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. Von Lilienfeld, "Quantum chemistry structures and properties of 134 kilo molecules," *Sci. Data* **1**, 140022 (2014).
- ⁹⁵M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. Von Lilienfeld, "Fast and accurate modeling of molecular atomization energies with machine learning," *Phys. Rev. Lett.* **108**, 058301 (2012).
- ⁹⁶J. S. Smith, O. Isayev, and A. E. Roitberg, "ANI-1, a data set of 20 million calculated off-equilibrium conformations for organic molecules," *Sci. Data* **4**, 170193 (2017).
- ⁹⁷F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
- ⁹⁸Y.-Y. Yang, S.-C. Lee, Y.-A. Chung, T.-E. Wu, S.-A. Chen, and H.-T. Lin, "libact: Pool-based active learning in python," Technical Report, National Taiwan University, 2017; [arXiv:1710.00379](https://arxiv.org/abs/1710.00379) [cs].
- ⁹⁹M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from <https://www.tensorflow.org>.
- ¹⁰⁰D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, edited by Y. Bengio and Y. LeCun (2015), [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- ¹⁰¹I. H. Witten, E. Frank, and M. A. Hall, *Data Mining: Practical Machine Learning Tools and Techniques*, Morgan Kaufmann Series in Data Management Systems (Morgan Kaufmann, 2017).